(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
30 May 2002 (30.05.2002)

PCT

(10) International Publication Number
WO 02/42959 A2

(51) International Patent Classification[7]: G06F 17/60

(21) International Application Number: PCT/EP01/13418

(22) International Filing Date:
16 November 2001 (16.11.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/718,256 22 November 2000 (22.11.2000) US

(71) Applicant: **KONINKLIJKE PHILIPS ELECTRONICS N.V.** [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).

(72) Inventor: **SCHAFFER, James, D.**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).

(74) Agent: **GROENENDAAL, Antonius, W., M.**; Internationaal Octrooibureau B.V., Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).

(81) Designated States *(national)*: CN, JP, KR, VN.

(84) Designated States *(regional)*: European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: TELEVISION PROGRAM RECOMMENDER WITH INTERVAL-BASED PROFILES FOR DETERMINING TIME-VARYING CONDITIONAL PROBABILITIES

(57) Abstract: A method and apparatus are disclosed for generating television program recommendations based on a time-windowed viewing history of a viewer. Changes in the viewing preferences are automatically identified. The viewing history profile is maintained as a series of viewing history windows. The length of each individual window, and the overall length of the entire viewing history itself (i.e., the number of such smaller viewing windows) can be fixed or varied. Insignificant or non-reoccurring features that can be deleted from the viewer profile without losing significant information are also identified. A feature can be deleted from the viewer profile when the corresponding frequency count for the feature fails to exceed a minimum threshold for a predefined period of time. The disclosed television programming recommender adapts the generated television program recommendations to changes in viewing preferences. The conditional probability of a feature (the likelihood that programs in a known class, e.g., the class of programs watched, will have the feature) is calculated as a function of time. For cyclical or periodic changes in viewing preferences, the disclosed television programming recommender adjusts the conditional probability of an attribute so that it more accurately reflects one or more similar earlier cycles or periods. For trends or permanent changes in viewing preferences, the television programming recommender generates television program recommendations using the most recent window(s) of the viewing history, which are more likely to reflect the trend (current viewing preferences), or extrapolates a value from an identified trend.

WO 02/42959 A2

Television program recommender with interval-based profiles for determining time-varying conditional probabilities

Field of the Invention

The present invention relates to television program recommenders, and more particularly, to a method and apparatus for generating television program recommendations.

Background of the Invention

As the number of channels available to television viewers has increased, along with the diversity of the programming content available on such channels, it has become increasingly challenging for television viewers to identify television programs of interest. Historically, television viewers identified television programs of interest by analyzing printed television program guides. Typically, such printed television program guides contained grids listing the available television programs by time and date, channel and title. As the number of television programs has increased, it has become increasingly difficult to effectively identify desirable television programs using such printed guides.

More recently, television program guides have become available in an electronic format, often referred to as electronic program guides (EPGs). Like printed television program guides, EPGs contain grids listing the available television programs by time and date, channel and title. Some EPGs, however, allow television viewers to sort or search the available television programs in accordance with personalized preferences. In addition, EPGs allow for on-screen presentation of the available television programs.

While EPGs allow viewers to identify desirable programs more efficiently than conventional printed guides, they suffer from a number of limitations, which if overcome, could further enhance the ability of viewers to identify desirable programs. For example, many viewers have a particular preference towards, or bias against, certain categories of programming, such as action-based programs or sports programming. Thus, the viewer preferences can be applied to the EPG to obtain a set of recommended programs that may be of interest to a particular viewer.

Thus, a number of tools have been proposed or suggested for recommending television programming. The Tivo™ system, for example, commercially available from Tivo, Inc., of Sunnyvale, California, allows viewers to rate shows using a "Thumbs Up and

Thumbs Down" feature and thereby indicate programs that the viewer likes and dislikes, respectively. Thereafter, the TiVo receiver matches the recorded viewer preferences with received program data, such as an EPG, to make recommendations tailored to each viewer. While such currently available television program recommenders help a user to identify television programs of interest, they suffer from a number of limitations, which if overcome, could greatly improve the performance in the generated recommendations.

For example, such television program recommenders typically record the viewer's prior viewing history in the form of a viewer profile, comprised of a number of program features and the corresponding number of occurrences of the feature in programs that were watched (and, optionally, not watched) by the viewer. The number of occurrences of a feature in programs that were watched (and, optionally, not watched) by the viewer is often referred to as the "feature count." Thus, as such conventional tools accumulate viewing histories over time, the number of features for which viewing history information is maintained progressively increases as new features are identified in programs that are watched by the viewer. Some identified features, however, may never again appear in another program watched by the viewer. Thus, unnecessary program features are often maintained in the viewer profile longer than necessary. There is currently no mechanism, however, that determines if a feature can be removed from the viewing history without loss of significant information.

Furthermore, with such television program recommenders, changes in the television viewing habits of a viewer, often referred to as "non-stationary viewing preferences," are not easily identified. In fact, the aggregate nature of the feature counts in the viewer profile will mask any temporal patterns in the viewing history.

A need therefore exists for methods and apparatus for generating television program recommendations that identify changing viewer preferences. A further need therefore exists for methods and apparatus for generating television program recommendations that adapt television program recommendations to changing viewer preferences. A further need exists for tools in such television program recommenders that maintain viewing histories in a more efficient manner.

Summary of the Invention

Generally, a method and apparatus are disclosed for generating television program recommendations based on the prior viewing history of a viewer. According to one aspect of the invention, changes in the viewing preferences are automatically identified.

According to another aspect of the invention, the viewing history is maintained as a series of viewing history windows, such as a series of consecutive periodic viewing history windows. The present invention also allows the length of each individual window, and the overall length of the entire viewing history itself (i.e., the number of such smaller viewing windows) to be varied.

Another aspect of the invention identifies insignificant or non-reoccurring features that can be deleted from the viewer profile without losing significant information. Generally, a feature (i.e., an attribute-value pair) can be deleted from the viewer profile when the corresponding frequency count for the feature fails to exceed a minimum threshold for a predefined period of time.

Once changing viewing preferences are identified, the disclosed television programming recommender can adapt the generated television program recommendations to such changes in viewing preferences. In particular, the present invention allows the conditional probability of a feature (the likelihood or probability that programs in the class of programs watched (or not watched) by the viewer will have the feature) to be calculated as a function of time. For cyclical or periodic changes in viewing preferences, such as repetitive changes corresponding to the time of day or time of year, the television programming recommender adjusts the conditional probability of an attribute so that it more accurately reflects one or more similar earlier cycles or periods. Likewise, for trends or permanent changes in viewing preferences, such as declining or increasing interest in a program, the television programming recommender could generate television program recommendations using the most recent window(s) of the viewing history, which are more likely to reflect the trend (current viewing preferences). Alternatively, television program recommendations could be generated using values extrapolated from a strong trend.

A more complete understanding of the present invention, as well as further features and advantages of the present invention, will be obtained by reference to the following detailed description and drawings.

Brief Description of the Drawings

FIG. 1 illustrates a television programming recommender in accordance with the present invention;

FIG. 2 illustrates the maintenance of a time-windowed viewing history in accordance with the present invention;

FIG. 3 is a sample table from the program database of FIG. 1;

FIG. 4 is a sample table from the viewer profile of FIG. 1;

FIG. 5 is a flow chart describing an exemplary viewer profile evaluation process embodying principles of the present invention;

FIG. 6 is a flow chart describing an exemplary viewer profile maintenance process embodying principles of the present invention;

FIGS. 7A and 7b, collectively, provide an example of the determination of the maximum likelihood estimates that are applied using linear trend testing and extrapolation techniques in accordance with the present invention; and

FIG. 8 is a flow chart describing an exemplary profile window maintenance routine incorporating features of the present invention.

Detailed Description

FIG. 1 illustrates a television programming recommender 100 in accordance with the present invention. As shown in FIG. 1, the television programming recommender 100 evaluates each of the programs in an electronic programming guide (EPG) 110 to identify programs of interest to a particular user. The set of recommended programs can be presented to the user using a set-top terminal/television 150, for example, using well known on-screen presentation techniques. The television programming recommender 100 identifies programs that a viewer might like, based on their prior viewing history 200, discussed further below in conjunction with FIG. 2.

According to one feature of the present invention, discussed further below in conjunction with FIG. 2, the television programming recommender 100 maintains the viewing history 200 as a series of viewing history windows that allow non-stationary viewing preferences to be identified. For example, the television programming recommender 100 can maintain the overall viewing history 200 as a series of consecutive periodic viewing history windows, such as one-month intervals. In one implementation, the length of each individual window, and the overall length of the entire viewing history itself (i.e., the number of such smaller viewing windows) may be varied. While the illustrative embodiments of the television programming recommender 100 discussed herein contemplate consecutive, adjacent, non-overlapping viewing history windows of equal duration, many variations are possible within the scope of the present invention, as would be apparent to a person of ordinary skill in the art.

As previously indicated, a viewer profile typically indicates the number of occurrences of each of a number of program attributes in programs that were watched or not

watched (or both) by the viewer. Thus, as conventional recommender tools accumulate viewing histories over time, the number of attributes for which viewing history information is maintained progressively increases. Another feature of the present invention identifies insignificant or non-reoccurring features and thereby controls the growth of viewer profiles. In one embodiment, the present invention removes a given attribute from the viewer profile when the corresponding frequency count for the feature fails to exceed a predefined minimum threshold, such as an estimated noise level, for a predefined period of time.

According to another feature of the present invention, once changing viewing preferences are identified, the television programming recommender 100 can adapt the generated television program recommendations to such changes in viewing preferences. In particular, the present invention allows the conditional probability of a feature (the likelihood that programs in the class of programs watched (or not watched) by the viewer will have the feature) to be calculated as a function of time.

For cyclical or periodic changes in viewing preferences, such as repetitive changes corresponding to the time of day or time of year, the television programming recommender 100 adjusts the conditional probability of a feature so that it corresponds more accurately to the expected value at the current point in its cycle. Among other benefits, the time-history approach to viewing histories associated with the present invention reduces the negative impact on programs that do not contain a feature that does not apply in a dissimilar later time interval (e.g., the off-season of a sport).

Likewise, for trends or permanent changes in viewing preferences, such as declining or increasing interest in a program, the television programming recommender 100 generates television program recommendations using the most recent window of the viewing history 200, which most likely reflects the trend (current viewing preferences), or adjusts the conditional probability to correspond to a value predicted by the trend.

For example, the viewing habits of many viewers may change temporarily during the fall season with the start of a new television season, or during the playoffs of a given sport. Likewise, the viewing habits of a given viewer may be different in the morning hours as compared to the evening hours. Thus, the present invention improves the performance of the television programming recommender 100 by detecting temporal patterns in the viewing history 200.

As shown in FIG. 1, the television programming recommender 100 contains a viewing history 200, a program database 300 and a viewer profile 400, each discussed further below in conjunction with FIGS. 2 through 4, respectively. Generally, the viewing history

200 identifies the programs that were watched (and/or not watched) by the viewer over a period of time. The program database 300 records information for each program that is available in a given time interval. The viewer profile 400 indicates the relative level of interest of the viewer in each corresponding program feature, based on viewing patterns.

In addition, the television programming recommender 100 includes a viewer profile evaluation process 500, a viewer profile maintenance process 600 and a profile window maintenance routine 800, each discussed further below in conjunction with FIGS. 5, 6 and 8, respectively. The viewer profile evaluation process 500 performs statistical analysis on the conditional probabilities in the viewer profile 400 and identifies temporal patterns in the viewing history (e.g., trends or periodic behavior). Once temporal patterns are identified by the viewer profile evaluation process 500, the corresponding conditional probability values may be appropriately adjusted to take non-stationary viewing preferences into account in accordance with the present invention. The viewer profile maintenance process 600 identifies and deletes insignificant features from the viewer profile 400. The profile window maintenance routine 800 provides an exemplary process for creating and deleting time-windows in the viewer profile 400, as appropriate, over time.

The television program recommender 100 may be embodied as any computing device, such as a personal computer or workstation, having a processor 120 and memory 130. In addition, the television programming recommender 100 may be embodied, for example, as the television program recommender described in United States Patent Application Serial No. 09/498,271, filed Feb. 4, 2000, entitled "Bayesian TV Show Recommender," (Attorney Docket No. 700690), as modified herein to carry out the features and functions of the present invention.

FIG. 2 conceptually illustrates the processing of the viewing history 200 in accordance with the present invention. As shown in FIG. 2, the viewing history 200 is maintained as a series of viewing windows, VH_N, that allow non-stationary viewing preferences to be identified. For example, the television programming recommender 100 can maintain the overall viewing history 200 as a series of consecutive periodic viewing windows, VH_N, such as one-month intervals. As shown in FIG. 2, the viewing history 200 can be maintained in a small buffer or memory location until the corresponding viewer profile 400 is generated or updated. In other words, the list of actual programs that were watched (or not watched) by the viewer only needs to be maintained until the corresponding conditional probabilities are determined for the viewer profile 400, as discussed below in conjunction with FIG. 4.

In one implementation, the length of each individual window, and the overall length of the entire viewing history itself (i.e., the number of such smaller viewing windows) may be varied. The viewing history windows, VH_1 and VH_K, shown in FIG. 2 correspond to the sets of shows watched by the viewer during time spans T_1 and T_K, respectively. The illustrative viewing history sub-sets, VH_1 and VH_K, correspond to time spans that are less than the entire time period covered by the viewing history 200. If a viewer changed his or her viewing preferences from time span T_1 to time span T_K, the programs corresponding to the new preferences would appear in time span T_K.

FIG. 3 is a sample table from the program database 300 of FIG. 1 that records information for each program that is available in a given time interval. As shown in FIG. 3, the program database 300 contains a plurality of records, such as records 305 through 320, each associated with a given program. For each program, the program database 300 indicates the date/time and channel associated with the program in fields 340 and 345, respectively. In addition, the title and genre for each program are identified in fields 350 and 355. Additional well-known attributes (not shown), such as actors, duration, and description of the program, can also be included in the program database 300.

In accordance with one feature of the present invention, the program database 300 can also record an indication of the recommendation score assigned to each program by the television programming recommender 100 in field 370. In this manner, the numerical scores can be displayed to the user in the electronic program guide sorted according to the numerical scores, and/or with each program directly or mapped onto a color spectrum or another visual cue that permits the user to quickly locate programs of interest.

FIG. 4 is a table illustrating an exemplary viewer profile 400. As shown in FIG. 4, the viewer profile 400 contains a plurality of records 405-413 each associated with a different program feature. In addition, for each feature set forth in column 440, the viewer profile 400 provides the corresponding number of occurrences of the feature in programs that were watched (and, optionally, not watched) by the viewer (positive and negative counts) for N different time windows. The positive and negative counts for each of N different time windows are set forth in columns 445 through 470, respectively. The positive and negative counts for the oldest time window 1 are set forth in columns 465 and 470, and the positive and negative counts for the most recent time window N are set forth in columns 445 and 450. Generally, when a new window (N + 1) is added to the viewer profile 400, the oldest window (1) shown in columns 465 and 470 "falls off" or can be stored in a historical viewing history database (not shown).

For example, if a given viewer watched a given sports program ten times on Channel 2 in the late afternoon in the most recent time window (N), then the positive counts associated with this features in the viewer profile 400 would be incremented by 10, and the negative counts would be 0 (zero). In the illustrative embodiment, the viewer profile 400 has N consecutive, adjacent, non-overlapping viewing windows of equal duration, such as one-month intervals. The overall length of the entire viewing history (i.e., the number, N, of such smaller viewing windows) may be empirically determined or programmatically established by the user. It is noted that the size of the window interval and the number of windows will impose limits on the nature of the temporal patterns that can effectively be identified. These parameters also interact with the rate at which a viewer watches television programs. FIG. 8, discussed below, provides an exemplary maintenance routine for the viewer profiles 400.

FIG. 5 is a flow chart describing an exemplary viewer profile evaluation process 500 embodying principles of the present invention. As previously indicated, the viewer profile evaluation process 500 performs statistical analysis on the conditional probabilities in the viewer profile 400 and identifies temporal patterns in the viewing history (e.g., trends or periodic behavior). Once temporal patterns are identified by the viewer profile evaluation process 500, the corresponding conditional probability values may be appropriately adjusted to take non-stationary viewing preferences into account in accordance with the present invention.

As shown in FIG. 5, the viewer profile evaluation process 500 initially retrieves the viewer profile 400 during step 510. Thereafter, the viewer profile evaluation process 500 performs a statistical analysis on the conditional probabilities for each feature during step 520. A test is performed during step 530 to determine if the statistical analysis indicates any trends or periodic viewing behavior. If it is determined during step 530 that there are no trends or periodic viewing behavior indicated in the viewer profile 400, then the conditional probabilities are computed during step 570 as the mean of the set of conditional probabilities, before program control terminates in step 580.

In other words, the conditional probabilities, CP, are computed during step 535 as follows:

$$CP = \frac{\#\ shows\ having\ attribute}{total\ \#\ shows\ in\ class}$$

If, however, it is determined during step 530 that there are trends or periodic viewing behavior indicated in the viewer profile 400, then a further test is performed during step 540 to distinguish between trends and periodic viewing behavior. If it is determined

during step 540 that the viewing behavior includes one or more trends, then the conditional probabilities for such corresponding features are computed during step 550 as the maximum likelihood estimate of the conditional probabilities given the trend. FIGS. 7A and 7B, discussed below, provide an example that illustrates the determination of the maximum likelihood estimates that are applied during step 550 (and also during step 560, discussed below) using linear trend testing and extrapolation techniques.

If, however, it is determined during step 540 that the viewing behavior includes one or more periodic tendencies, then the conditional probability for such corresponding features are computed during step 560 as the maximum likelihood estimate of the conditional probabilities given the periodic model (to more closely resemble those of corresponding earlier time periods). For example, if a particular viewer tends to watch a lot of baseball during the World Series, then the corresponding conditional probabilities should be adjusted in an upward manner for each time interval when the World Series is on. Likewise, the corresponding conditional probabilities should be adjusted in a downward manner for each time interval when the World Series is not on to thereby reduce the negative impact on programs that do not contain such corresponding features in a dissimilar later time interval (e.g., the off-season for baseball). Program control terminates during step 580.

FIG. 6 is a flow chart describing an exemplary viewer profile maintenance process 600 embodying principles of the present invention. As previously indicated, the viewer profile maintenance process 600 identifies and deletes insignificant attributes from the viewer profile 400. As shown in FIG. 6, the viewer profile maintenance process 600 initially retrieves the viewer profile 400 during step 610.

A test is performed during step 620 to determine if the profile 400 contains any features having a frequency count that falls below a predefined or dynamically determined noise threshold for a predefined consecutive number of windows. The noise threshold may be determined, for example, by analyzing the frequency counts of the randomly sampled negative examples. If it is determined during step 620 that one or more features have a frequency count that falls below a predefined noise threshold for a predefined consecutive number of windows, then such insignificant features are deleted from the viewer profile 400 during step 630.

For example, as shown in FIG. 4, the feature "early afternoon programs" represents a feature that was identified once in the random sampling of programs not watched by the viewer. If the window size of the viewer profile 400 is one-month and a total of N months are maintained in the viewing history, then the oldest month will be deleted and the

feature "early afternoon programs" can also be deleted for all months in the viewer profile 400 as not contributing to the current interests of the viewer. Program control terminates during step 640.

Linear Trend Testing And Extrapolation Techniques

FIGS. 7A and 7B, collectively, provide an example of the determination of the maximum likelihood estimates that are applied during steps 550 and 560, using linear trend testing and extrapolation techniques. FIG. 7A shows a portion of an exemplary viewer profile 700. As shown in FIG. 7A, the viewer profile 700 contains a plurality of records 705-707 indicating the total number of programs, feature counts for a representative feature and corresponding estimated conditional probabilities, respectively, for five different time windows. It is noted that the estimated conditional probabilities set forth in record 707 are computed only for the positive examples. In the illustrative example, however, there is no trend for the negative examples, and the estimated conditional probabilities would be zero.

FIG. 7B illustrates a regression line 770 that provides a best-fit of the five samples shown in FIG. 7A. In addition, FIG. 7B illustrates the prediction of the sixth (or any future) sample 780 using an extrapolation of the linear trend. As shown in FIG. 7B, a linear regression analysis is performed on the sample data 761-765 using well-known techniques to identify the regression line 770 having the lowest sum-square-error. Once the regression line 770 is obtained it can be extrapolated using known techniques to obtain the conditional probability 780 for the next (and any subsequent) interval. Generally, the regression line 770 is identified by its slope (m) and intersection (b), in a known manner. Once m and b are obtained, any future value, y can be predicted as follows:.

$$y = mx + b.$$

It is noted that since we are dealing with probability values, the extrapolation performed to predict a future sample 780 must be constrained between a range of zero (0) and one (1). While the example described above in conjunction with FIGS. 7A and 7B illustrate the determination of maximum likelihood estimates in the presence of linear trends, the same basic concepts can be applied to other trends as well, such as oscillations (cyclic behavior), as would be apparent to a person of ordinary skill in the art. It is further noted that combinations of the linear trends or cyclical/periodic viewing behavior described in the illustrative embodiment herein, such as a rising oscillation, are also possible. Although significant amounts of data are required to characterize such behavior, the techniques described herein can be applied to these more complex models, as would be apparent to a person of ordinary skill in the art, based on the disclosure herein.

FIG. 8 is a flow chart describing an exemplary profile window maintenance routine 800 for the viewer profiles 400. As shown in FIG. 8, the profile window maintenance routine 800 initially detects programs watched by the viewer during step 810. Thereafter, the profile window maintenance routine 800 performs a test during step 820 to determine if the current time is within the current window. In other words, the test performed during step 820 determines if it is time to create a new window in the profile 400.

If it is determined during step 820 that the current time is not within the current window, then the current window pointer is incremented during step 830, all frequency counts in the current window are set to zero during step 840, and the viewer profile maintenance process 600, discussed above in conjunction with FIG. 6, is executed during step 850 (to delete insignificant features). Program control then proceeds to step 860

If, however, it is determined during step 820 that the current time is within the current window, then the frequency counts in the current window of all features associated with the program are incremented during step 860. Thereafter, program control terminates. The negative counts would be processed in a similar manner.

It is to be understood that the embodiments and variations shown and described herein are merely illustrative of the principles of this invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention. For example, while the present invention has been illustrated in the context of a television program recommender, the invention could be applied to other sorts of user profiles, such as those based on purchase histories. Generally, the present invention can be applied to any profiles that monitor repeatable behavior.

CLAIMS:

1. A method for managing the storage of a user profile (400), comprising:

observing behavior of said user over time; and

maintaining said user profile (400) as a plurality of viewing history windows (VH_K), wherein each of said viewing history windows (VH_K) corresponds to a different time interval.

2. The method of claim 1, wherein said user profile (400) is associated with a television program recommender (100).

3. The method of claim 1, wherein said behavior is a set of programs that have been watched by a user.

4. The method of claim 1, wherein said behavior is a set of purchases by a user.

5. The method of claim 1, wherein said behavior indicates content that was liked and/or content that was disliked by a user.

6. The method of claim 1, wherein each of said plurality of viewing history windows (VH_K) can be of variable length.

7. The method of claim 1, wherein each of said plurality of viewing history windows (VH_K) can be varied for a given program feature.

8. The method of claim 1 further comprising the steps of:

obtaining a behavior history (200) indicating a number of occurrences of a plurality of predefined features of behavior performed by said user for a plurality of different time intervals; and

deleting from said user profile (400) any of said features having a number of occurrences that falls below a threshold for a predefined number of said time intervals.

9. The method of claim 8, wherein said threshold is dynamically determined.

10. The method of claim 8, wherein said user profile (400) further includes a number of occurrences of a plurality of features in at least a portion of programs that were not watched by said viewer for each of said plurality of different time intervals.

11. The method of claim 10, wherein said threshold is obtained by analyzing said number of occurrences of said plurality of features in at least a portion of the programs that were not watched by said viewer for each of said plurality of different time intervals.

12. The method of claim 8 further comprising the step of:

analyzing said number of occurrences for said features to identify a trend in said viewing behavior.

13. The method of claim 12, further comprising the step of adjusting a conditional probability of any feature exhibiting a trend using said number of occurrences for more recent time intervals.

14. The method of claim 12, further comprising the step of extrapolating conditional probability values of any feature exhibiting a trend to estimate a conditional probability value for a future period of time.

15. The method of claim 12 further comprising the step of:

analyzing said number of occurrences for said features to identify substantially periodic behavior in said viewing behavior.

16. The method of claim 15, further comprising the step of adjusting a conditional probability of any feature exhibiting periodic behavior using said number of occurrences for an earlier similar time interval.

17. A system for managing the storage of a viewer profile (400) in a television program recommender (100), comprising:

a memory (130) for storing computer readable code; and

a processor (120) operatively coupled to said memory (130), said processor (120) configured to:

obtain a viewing history (200) indicating a set of programs that have been watched by a user; and

maintain said viewing profile (400) as a plurality of viewing history windows (VH_K), wherein each of said viewing history windows (VH_K) corresponds to a different time interval.

18. The system of claim 17 said processor (120) further configured to:

obtain a viewing history (200) indicating a number of occurrences of a plurality of features in programs that were watched by a viewer for a plurality of different time intervals; and

delete from said viewer profile (400) any of said features having a number of occurrences that falls below a threshold for a predefined number of said time intervals.

19. The system of claim 18, wherein said viewer profile (400) further includes a number of occurrences of a plurality of features in at least a portion of the programs that were not watched by said viewer for each of said plurality of different time intervals.

20. The system of claim 19, wherein said threshold is obtained by analyzing said number of occurrences of said plurality of features in at least a portion of the programs that were not watched by said viewer for each of said plurality of different time intervals.

21. The system of claim 18 said processor (120) further configured to:

analyze said number of occurrences for said features to identify a trend in said viewing behavior.

22. The system of claim 21, wherein said processor (120) is further configured to extrapolate conditional probability values of any feature exhibiting a trend to estimate a conditional probability value for a future period of time.

23. The system of claim 21 said processor (120) further configured to:

analyze said number of occurrences for said features to identify substantially periodic behavior in said viewing behavior.

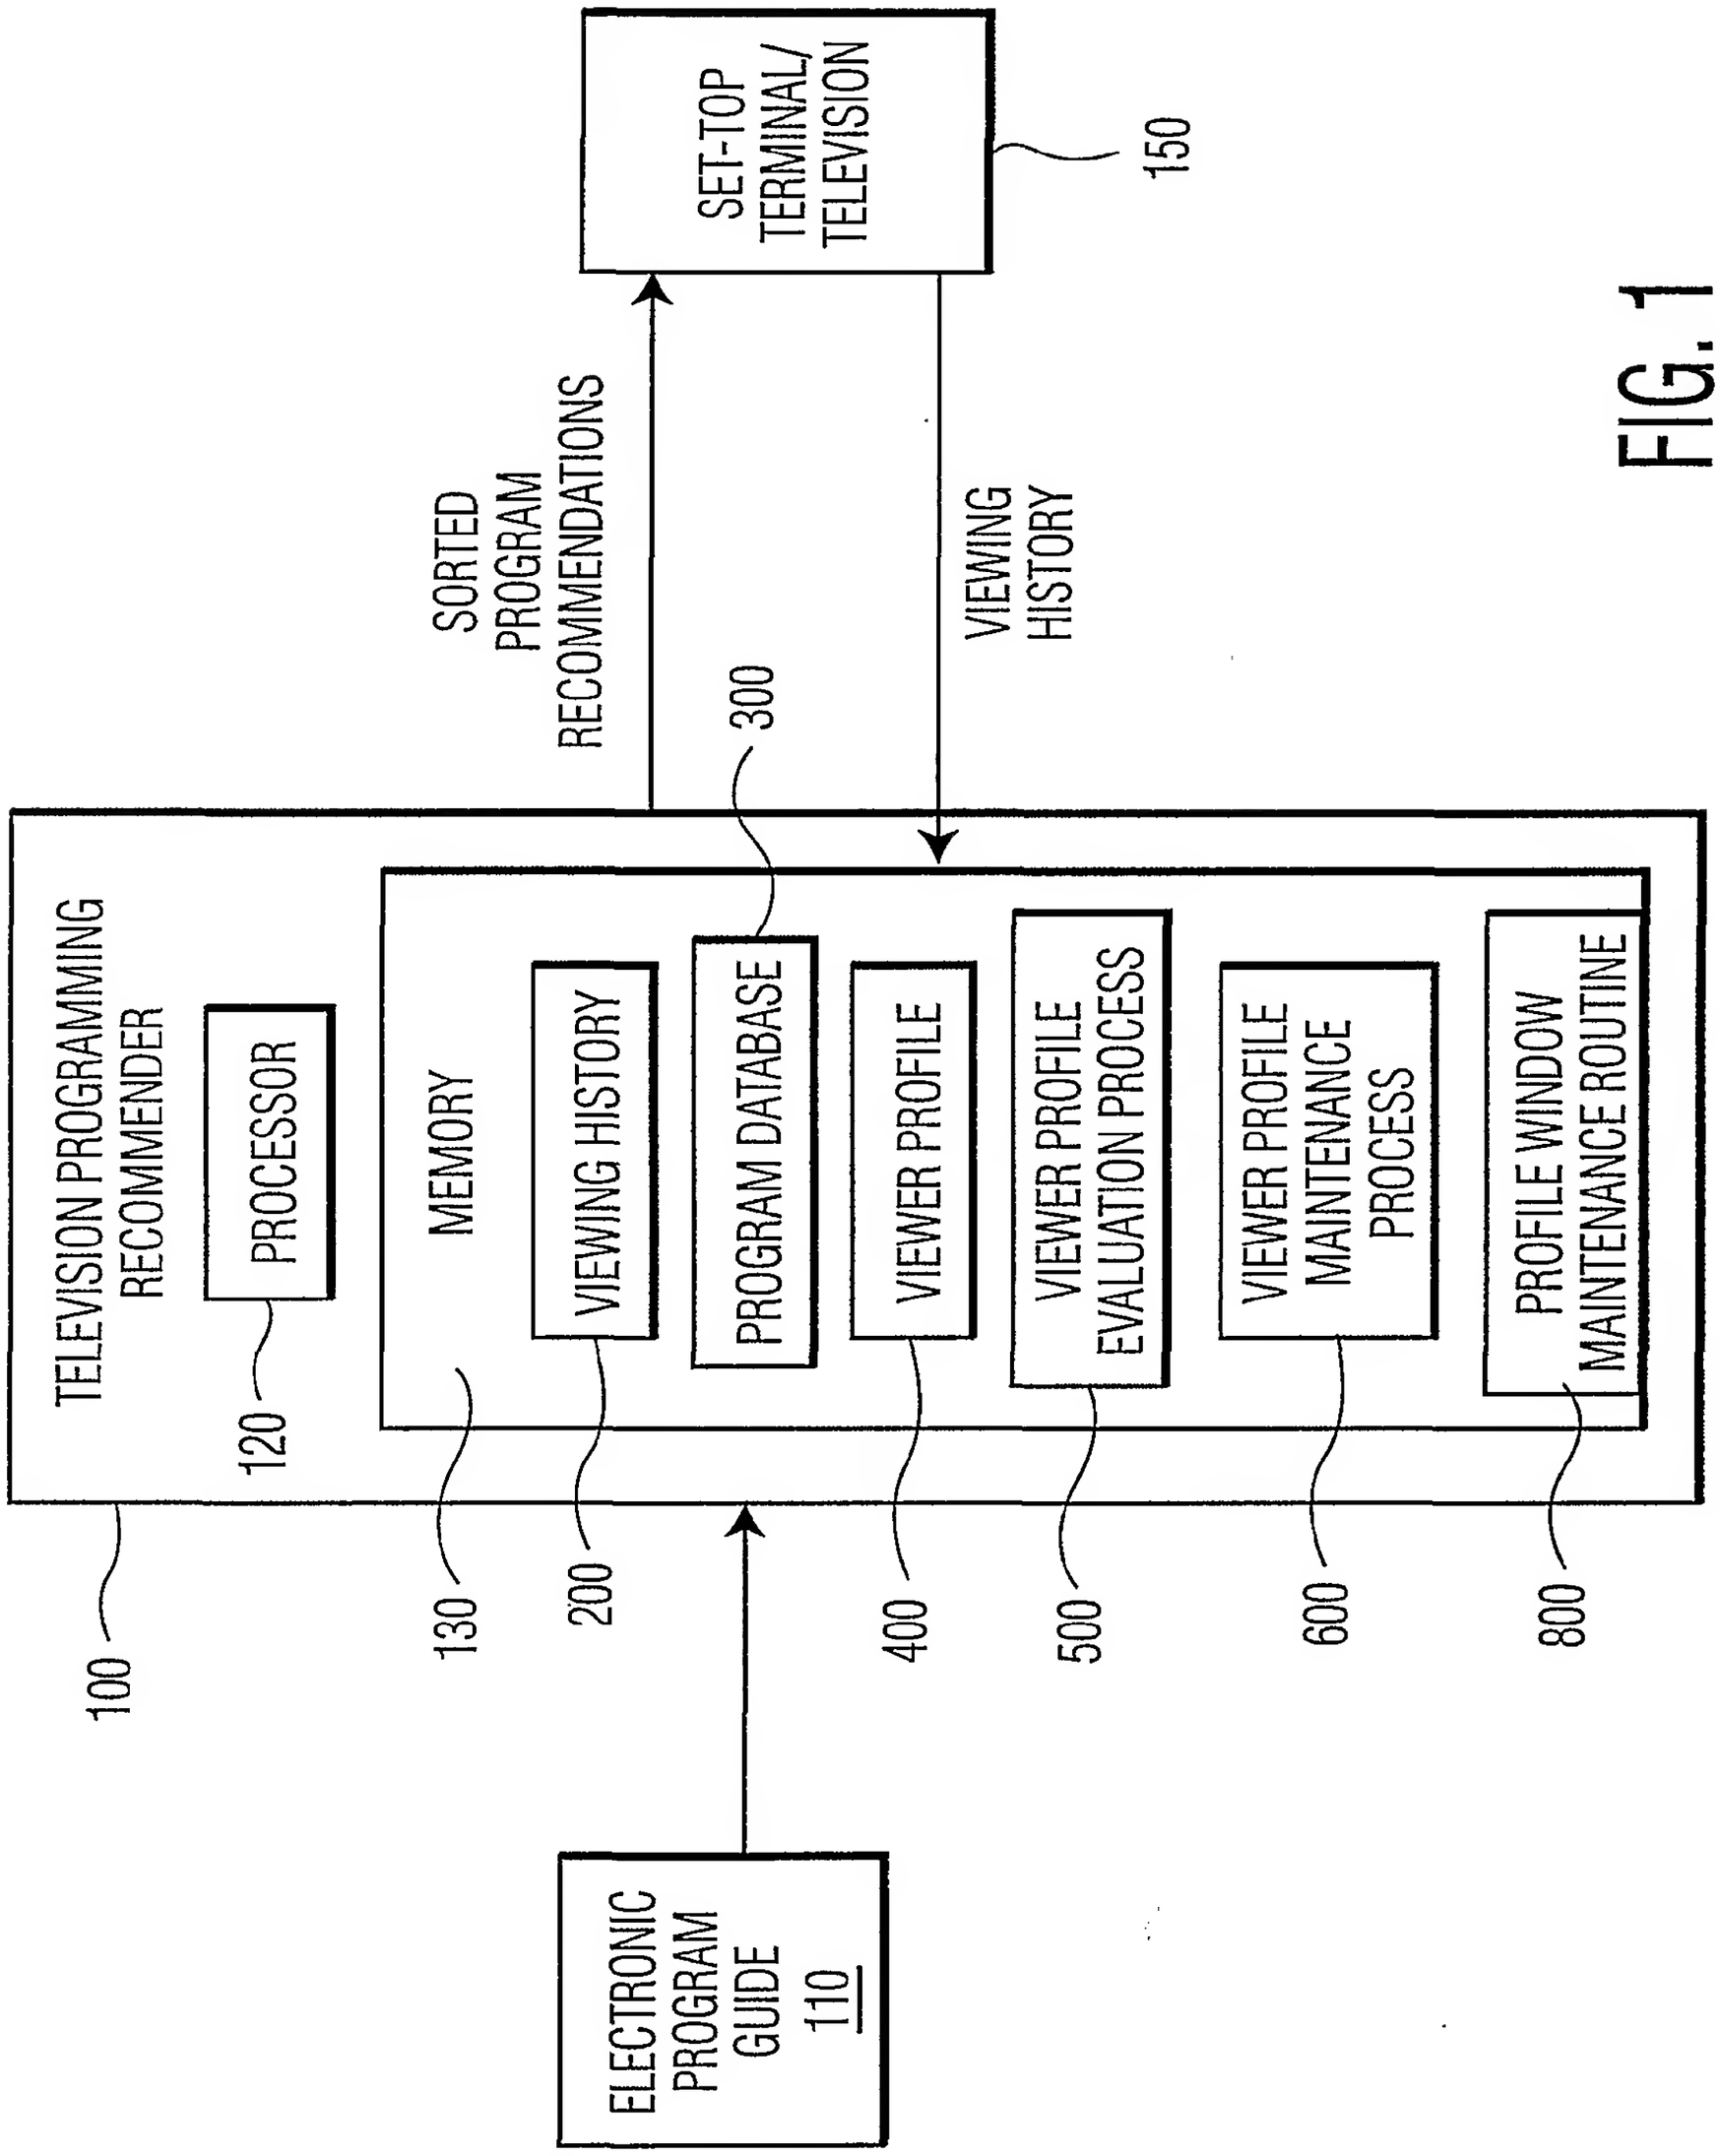
SET-TOP TERMINAL/ TELEVISION
150
SORTED PROGRAM RECOMMENDATIONS
VIEWING HISTORY
300
TELEVISION PROGRAMMING RECOMMENDER
PROCESSOR
120
MEMORY
VIEWING HISTORY
PROGRAM DATABASE
VIEWER PROFILE
VIEWER PROFILE EVALUATION PROCESS
VIEWER PROFILE MAINTENANCE PROCESS
PROFILE WINDOW MAINTENANCE ROUTINE
100
130
200
400
500
600
800
ELECTRONIC PROGRAM GUIDE
110

FIG. 1

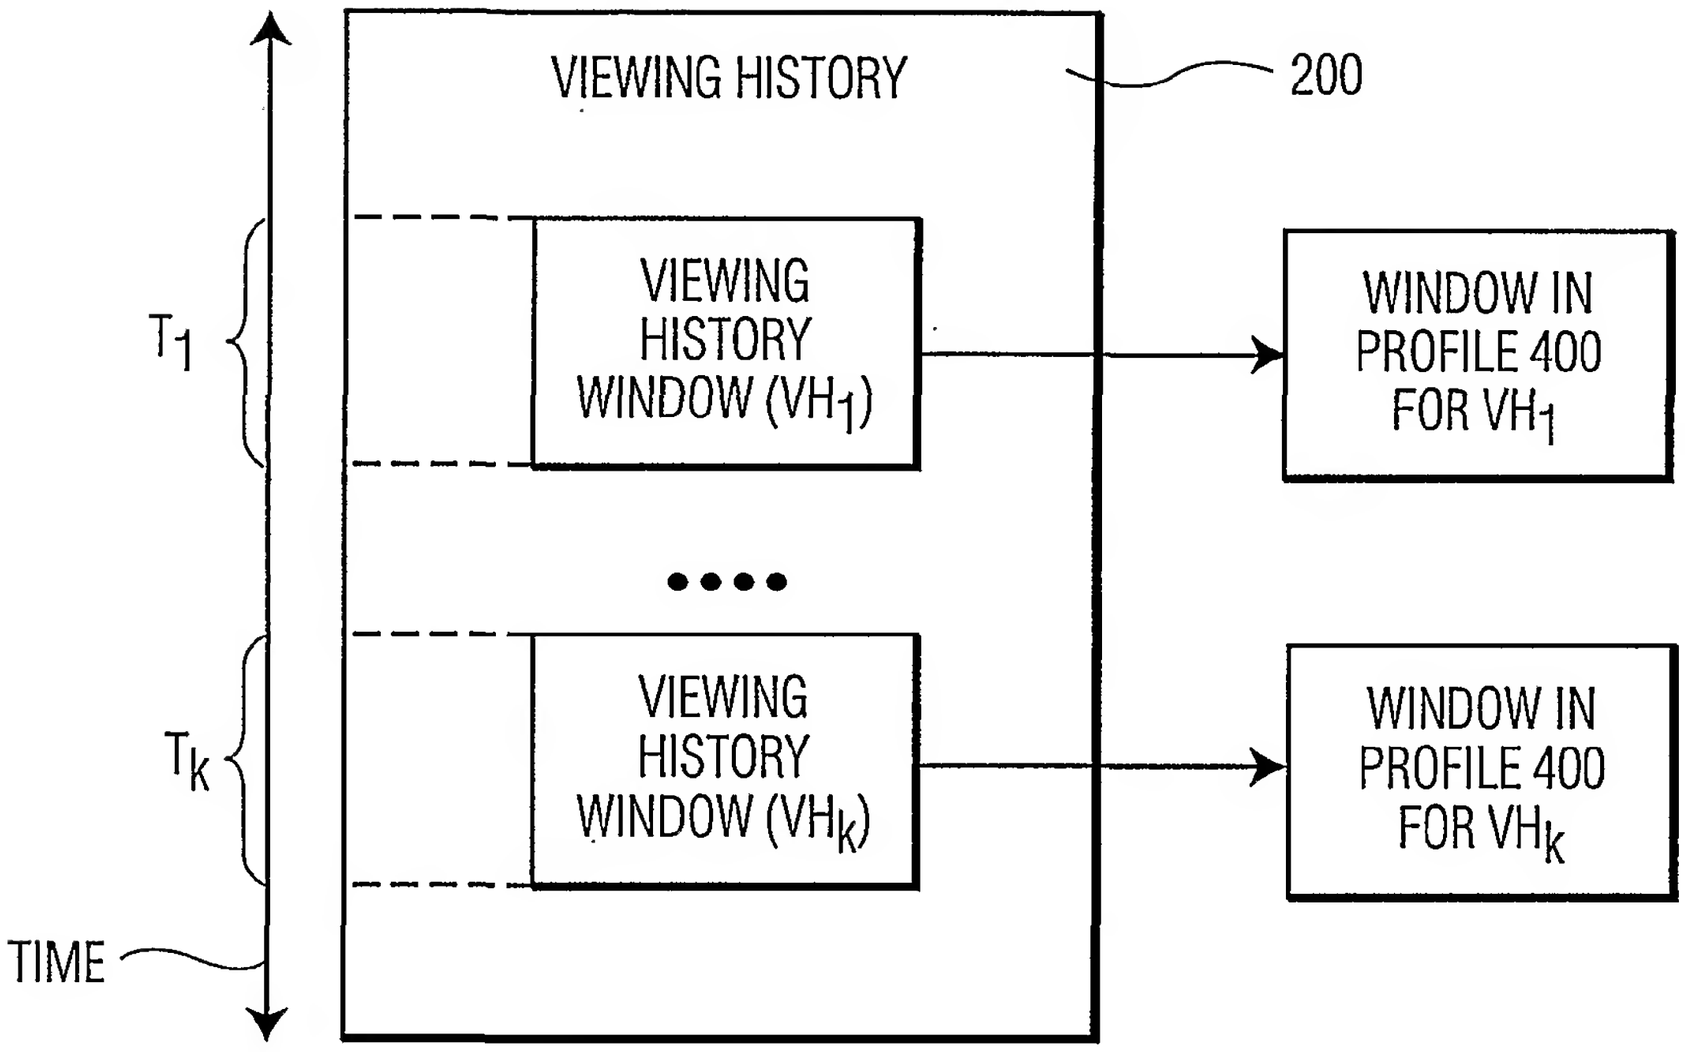

FIG. 2

PROGRAM DATABASE 300

DATE/TIME 340	CHANNEL 345	TITLE 350	GENRE 355	•••	RECOMMENDER SCORE 370
11/18/99 -- 8:00 P.M.	CH 1	LUCY	COMEDY		55
11/18/99 -- 8:30 P.M.	CH 1	AL'S FAMILY	SITCOM		78
•••					
11/18/99 -- 9:00 P.M.	CH 3	YOUR HOUSE	DRAMA		96

305
310
320

FIG. 3

VIEWER PROFILE ~ 400

		WINDOW N		WINDOW N-1			WINDOW 1	
	ATTRIBUTE 440	POSITIVE COUNTS 445	NEGATIVE COUNTS 450	POSITIVE COUNTS 455	NEGATIVE COUNTS 460	•••	POSITIVE COUNTS 465	NEGATIVE COUNTS 470
405	CHANNEL 2	10	0	8	0		6	0
406	CHANNEL 4	3	1	2	0		0	2
407	CHANNEL 7	4	4	3	2		5	1
	•••							
409	SPORTS CHANNEL	10	0	6	4		4	5
•••	MUSIC CHANNEL	1	0	0	1		1	1
	•••							
410	MORNING PROGRAMS	2	2	3	1		2	4
411	EARLY AFTERNOON PROGRAMS	0	0	0	0		0	1
412	LATE AFTERNOON PROGRAMS	10	0	6	4		3	2
413	EVENING PROGRAMS	6	4	3	5		4	3
	•••							

FIG. 4

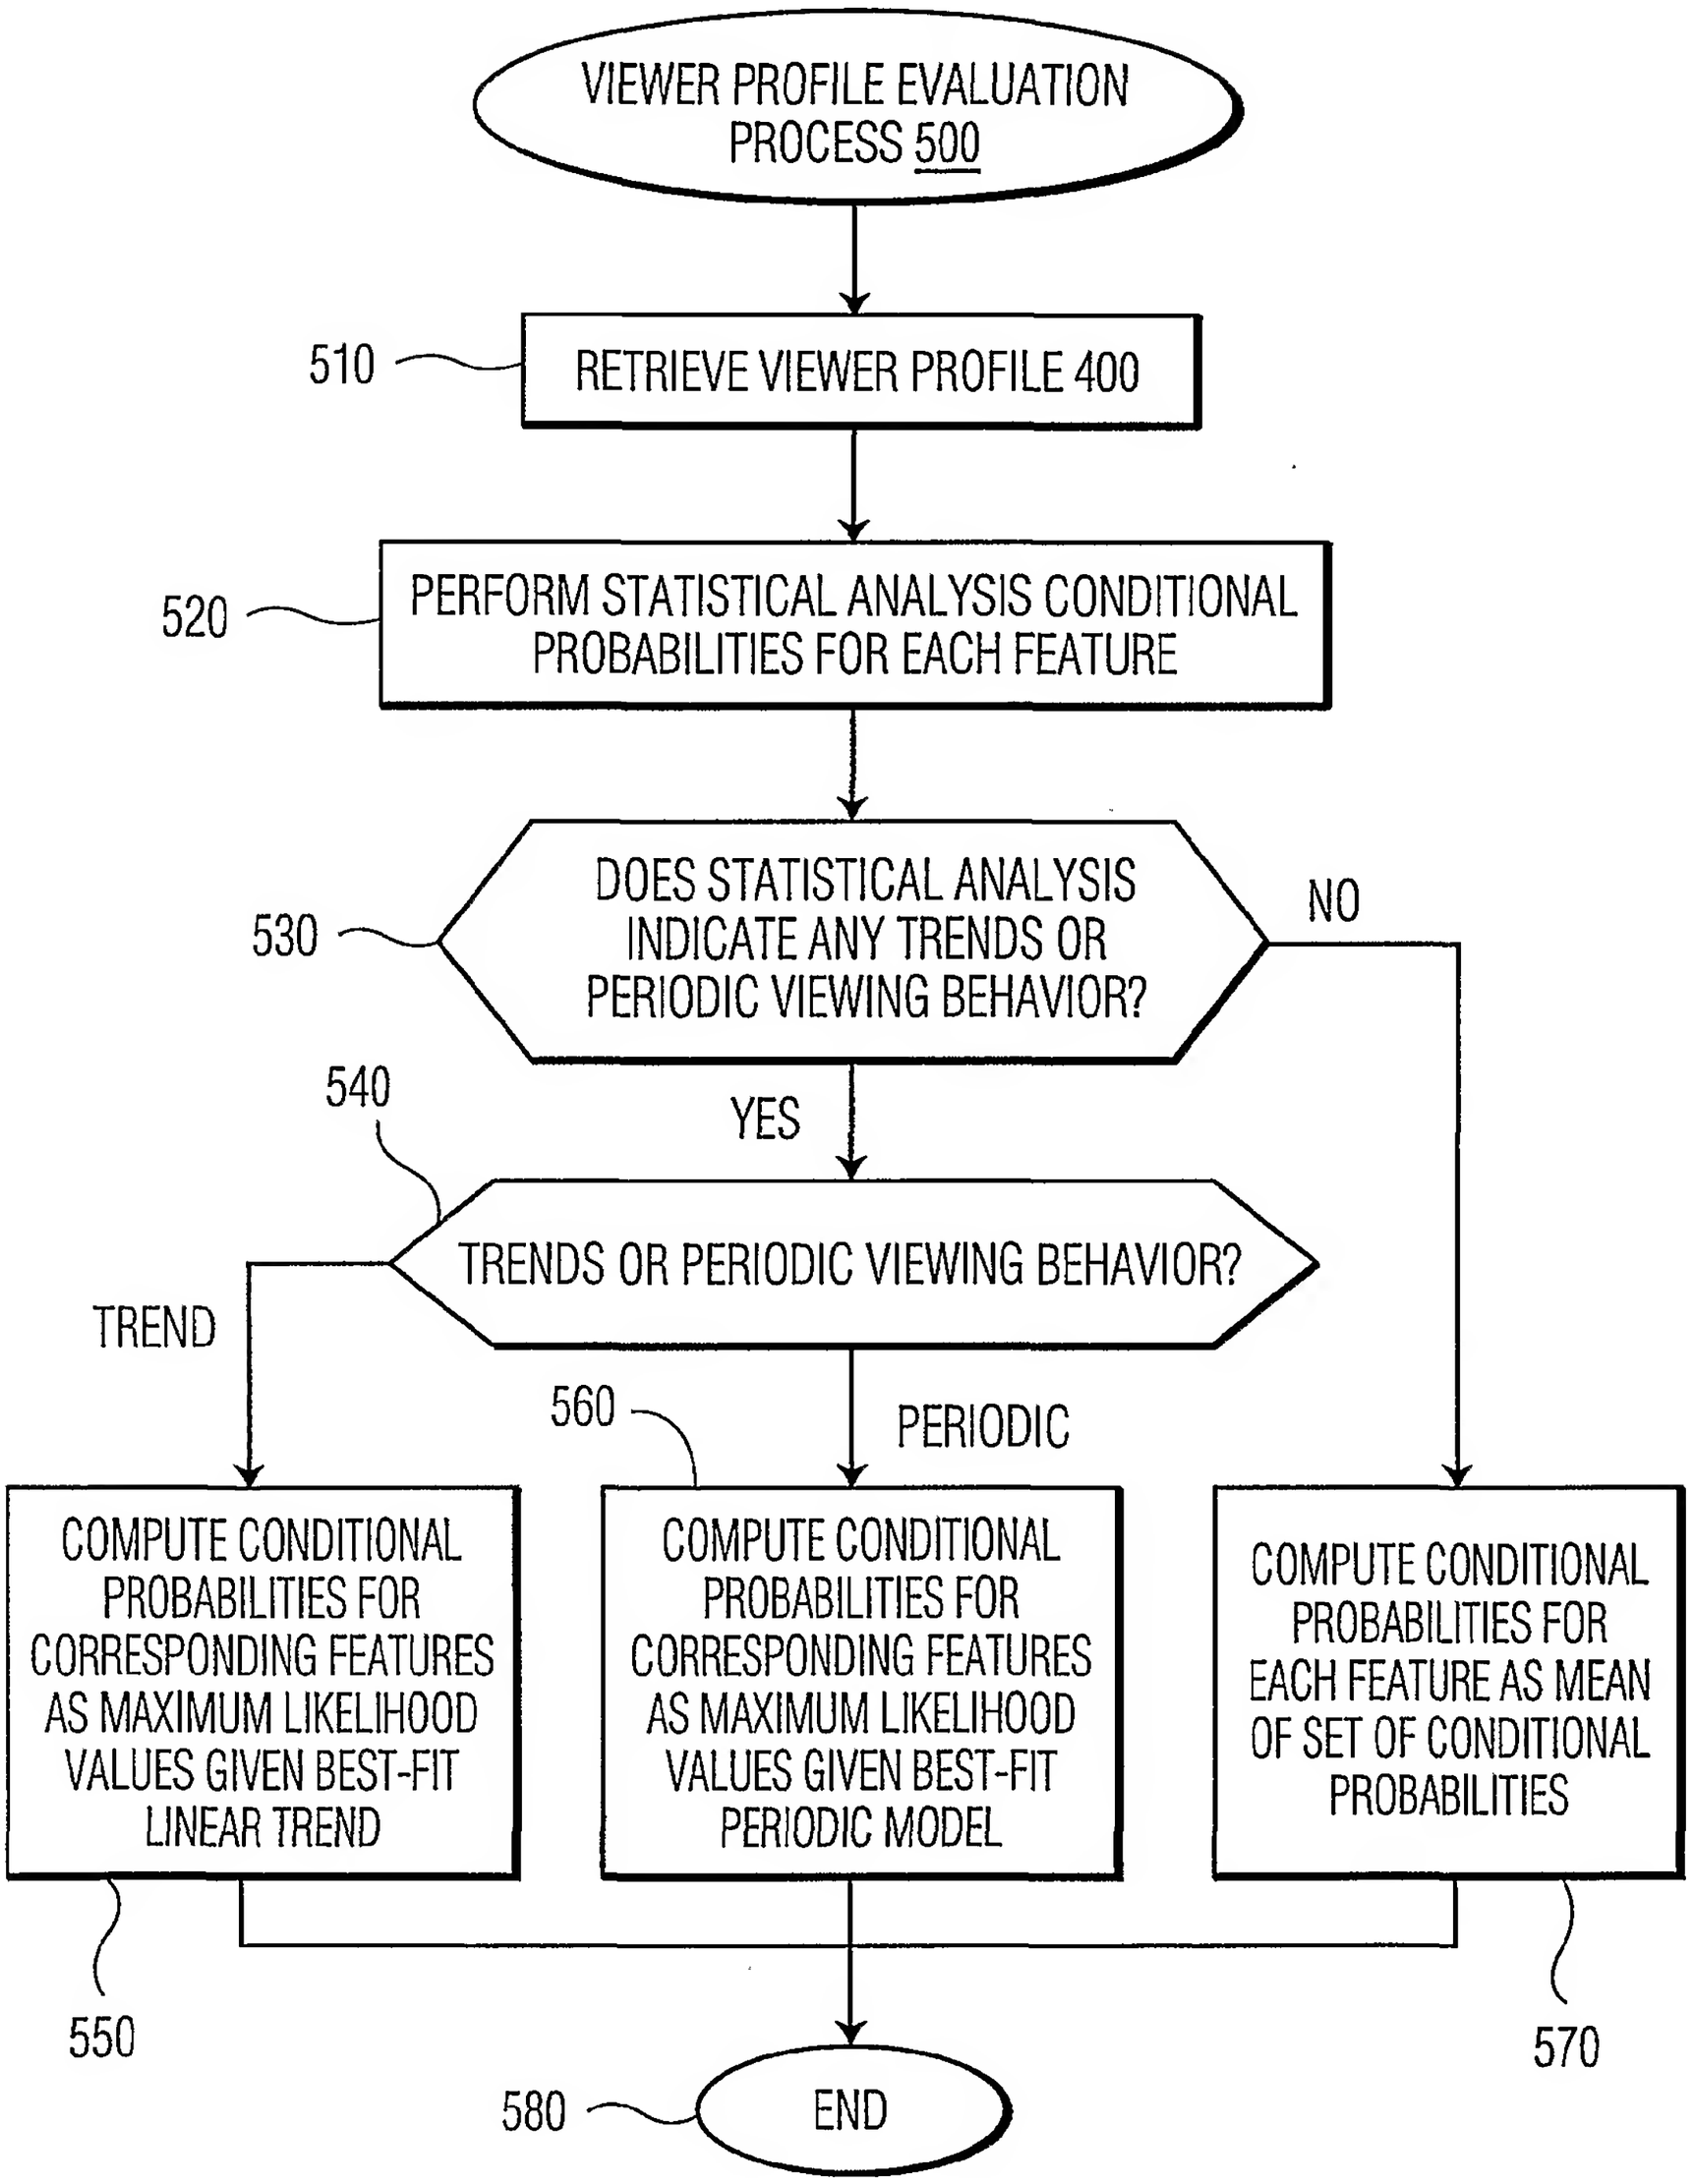
VIEWER PROFILE EVALUATION PROCESS 500
510
RETRIEVE VIEWER PROFILE 400
520
PERFORM STATISTICAL ANALYSIS CONDITIONAL PROBABILITIES FOR EACH FEATURE
530
DOES STATISTICAL ANALYSIS INDICATE ANY TRENDS OR PERIODIC VIEWING BEHAVIOR?
NO
YES
540
TRENDS OR PERIODIC VIEWING BEHAVIOR?
TREND
560
PERIODIC
COMPUTE CONDITIONAL PROBABILITIES FOR CORRESPONDING FEATURES AS MAXIMUM LIKELIHOOD VALUES GIVEN BEST-FIT LINEAR TREND
COMPUTE CONDITIONAL PROBABILITIES FOR CORRESPONDING FEATURES AS MAXIMUM LIKELIHOOD VALUES GIVEN BEST-FIT PERIODIC MODEL
COMPUTE CONDITIONAL PROBABILITIES FOR EACH FEATURE AS MEAN OF SET OF CONDITIONAL PROBABILITIES
550
570
580
END

FIG. 5

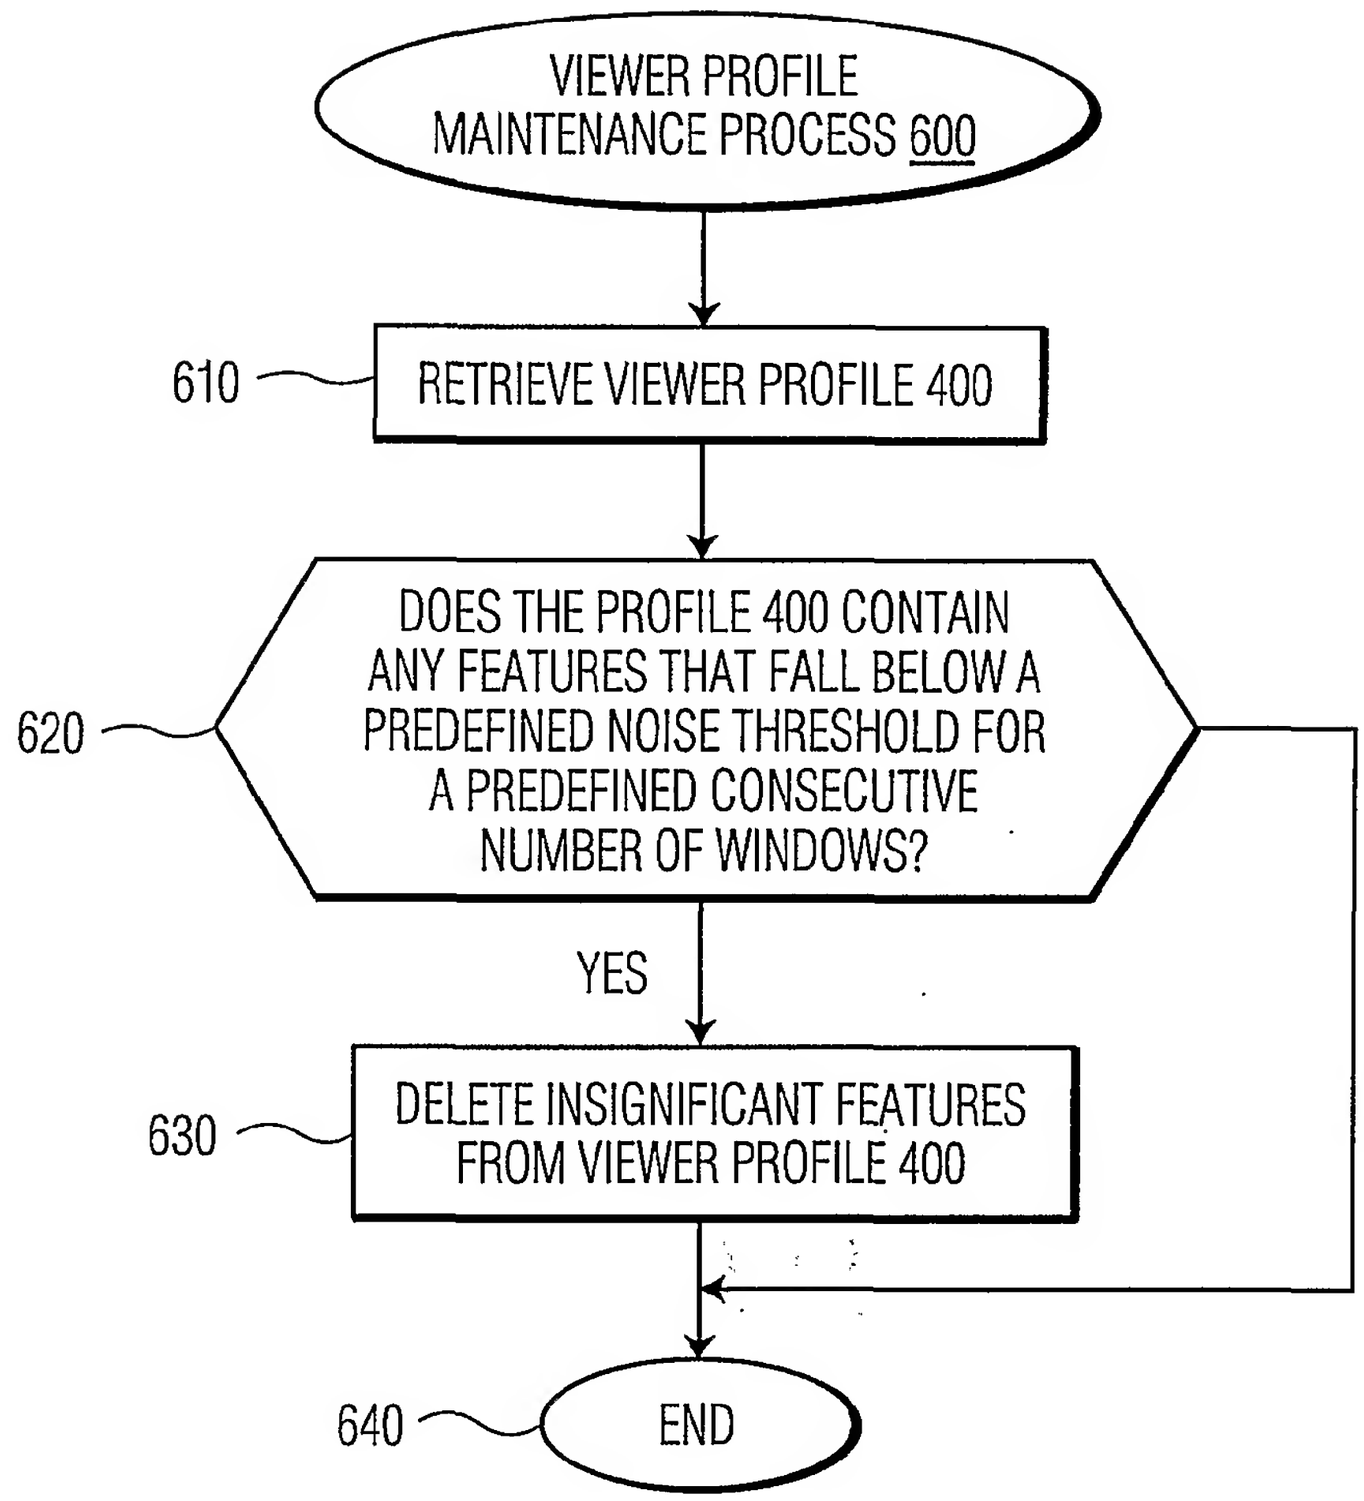
VIEWER PROFILE MAINTENANCE PROCESS 600
610
RETRIEVE VIEWER PROFILE 400
620
DOES THE PROFILE 400 CONTAIN ANY FEATURES THAT FALL BELOW A PREDEFINED NOISE THRESHOLD FOR A PREDEFINED CONSECUTIVE NUMBER OF WINDOWS?
YES
630
DELETE INSIGNIFICANT FEATURES FROM VIEWER PROFILE 400
640
END

FIG. 6

SAMPLE PROFILE — 700

		WINDOW 5		WINDOW 4		WINDOW 3		WINDOW 2		WINDOW 1	
		POS. COUNTS 710	NEG. COUNTS 715	POS. COUNTS 720	NEG. COUNTS 725	POS. COUNTS 730	NEG. COUNTS 735	POS. COUNTS 740	NEG. COUNTS 745	POS. COUNTS 750	NEG. COUNTS 755
705	TOTAL PROGRAMS	42	42	40	40	45	45	28	28	24	24
706	FEATURE COUNT	4	0	3	0	3	0	1	0	1	0
707	ESTIMATED COND'L PROBABILITY	0.0952	--	0.075	--	0.067	--	0.0357	--	0.0417	--

FIG. 7A

FIG. 7B

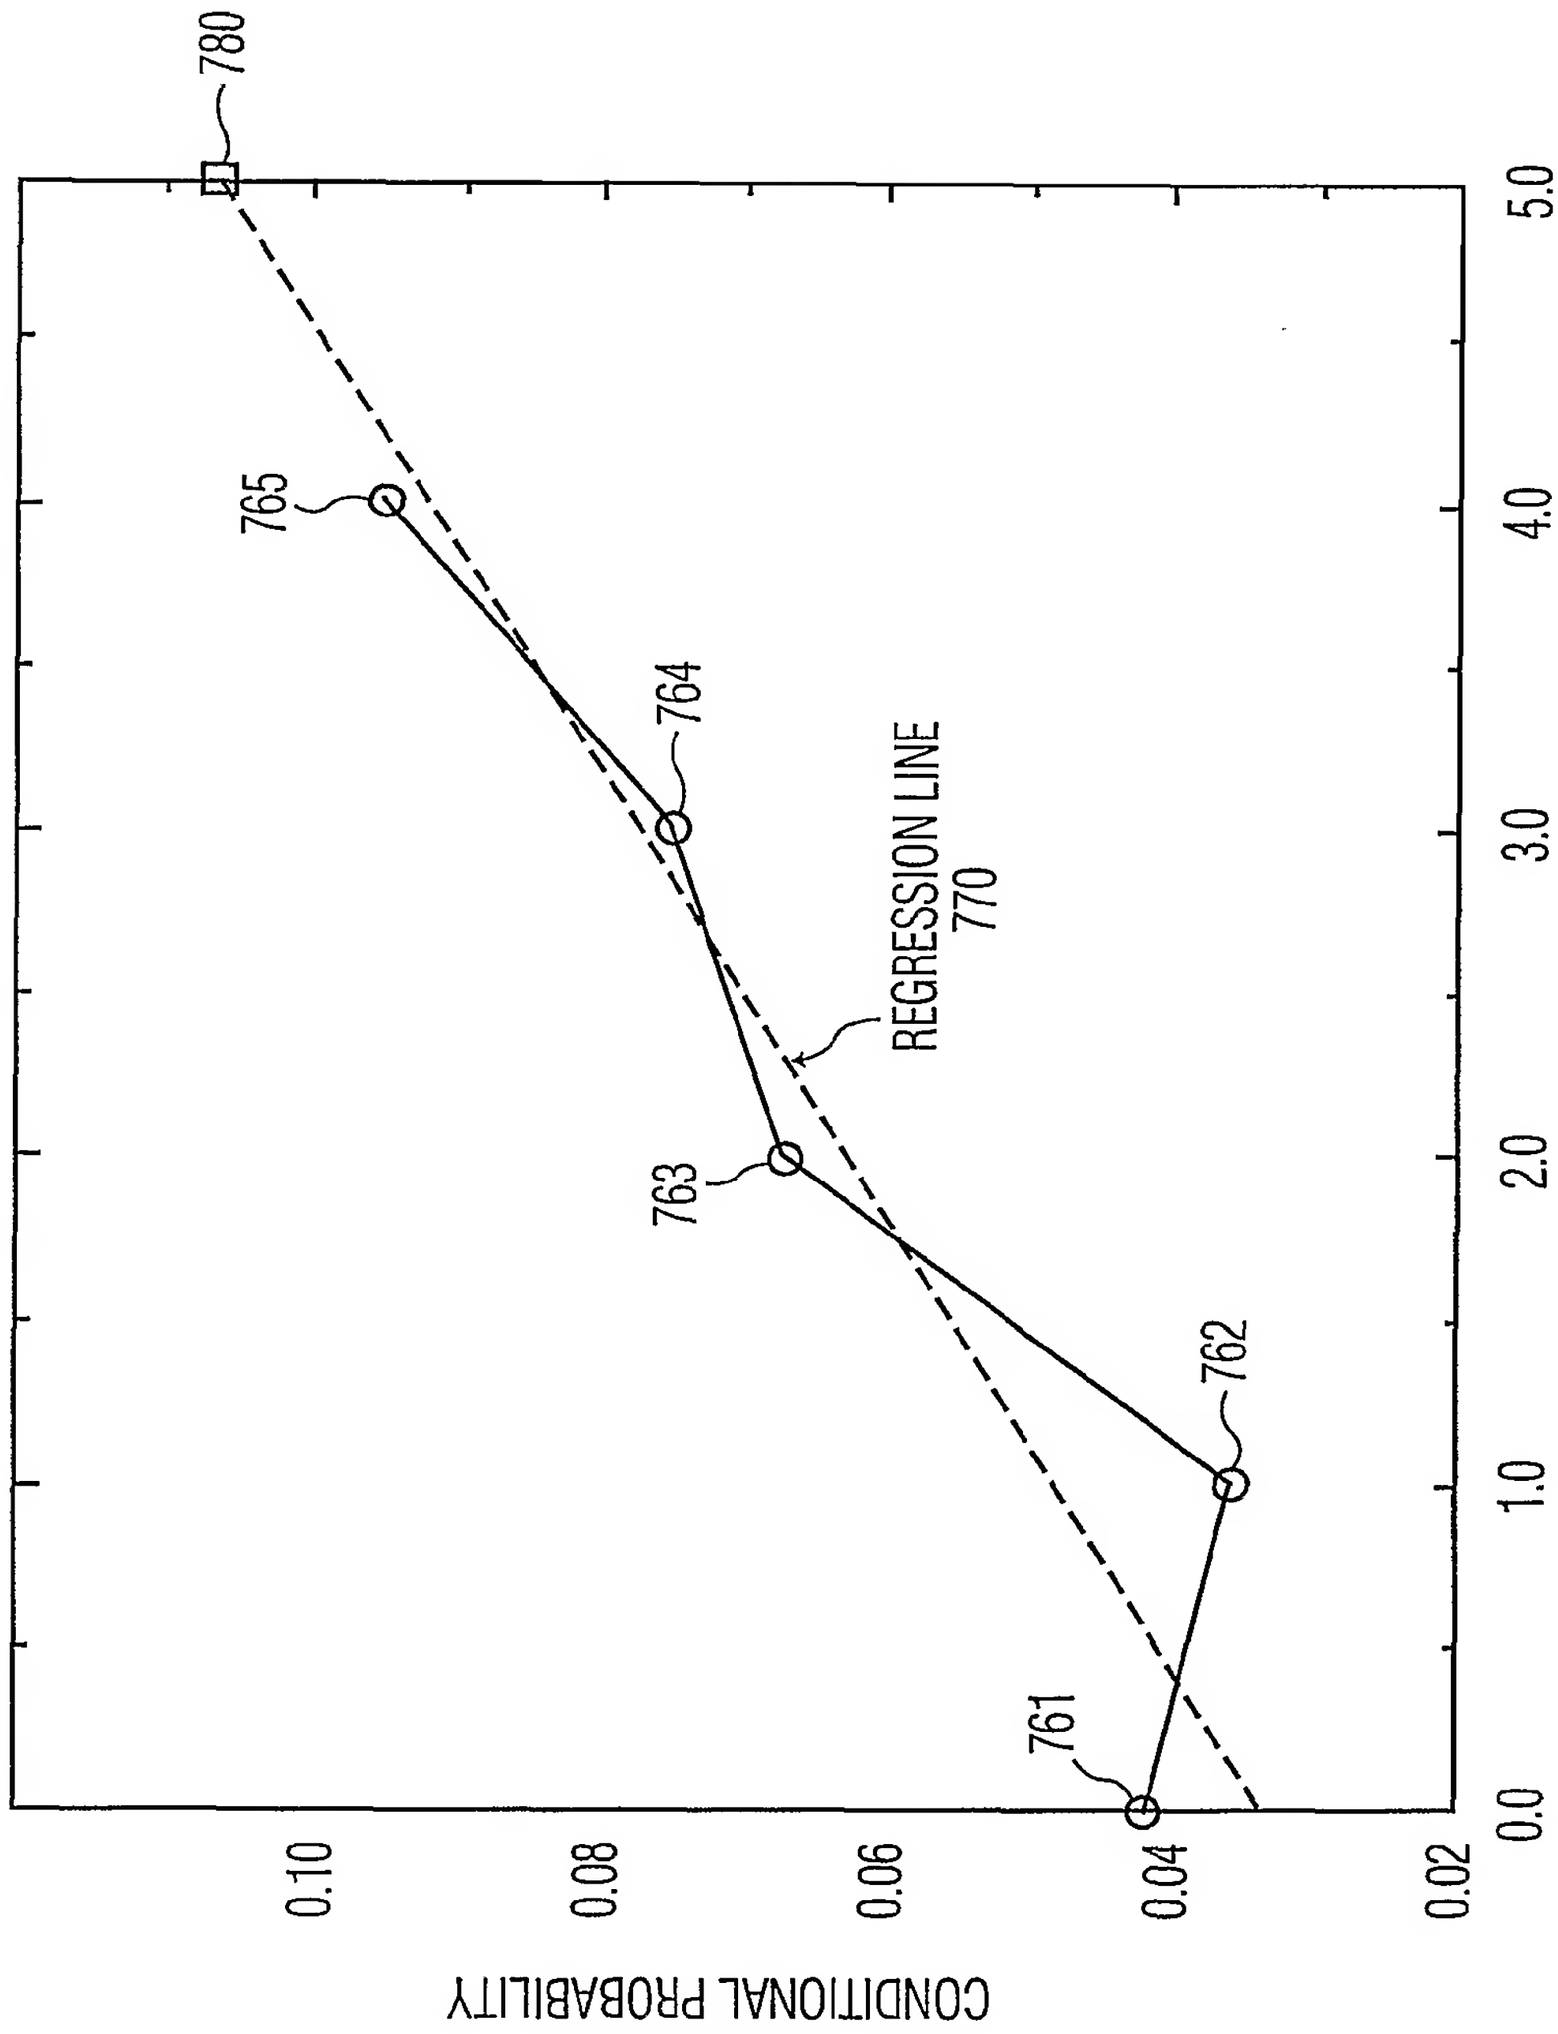
780
765
764
763
762
761
REGRESSION LINE
770
CONDITIONAL PROBABILITY
0.10
0.08
0.06
0.04
0.02
0.0
1.0
2.0
3.0
4.0
5.0
TEMPORAL SAMPLE

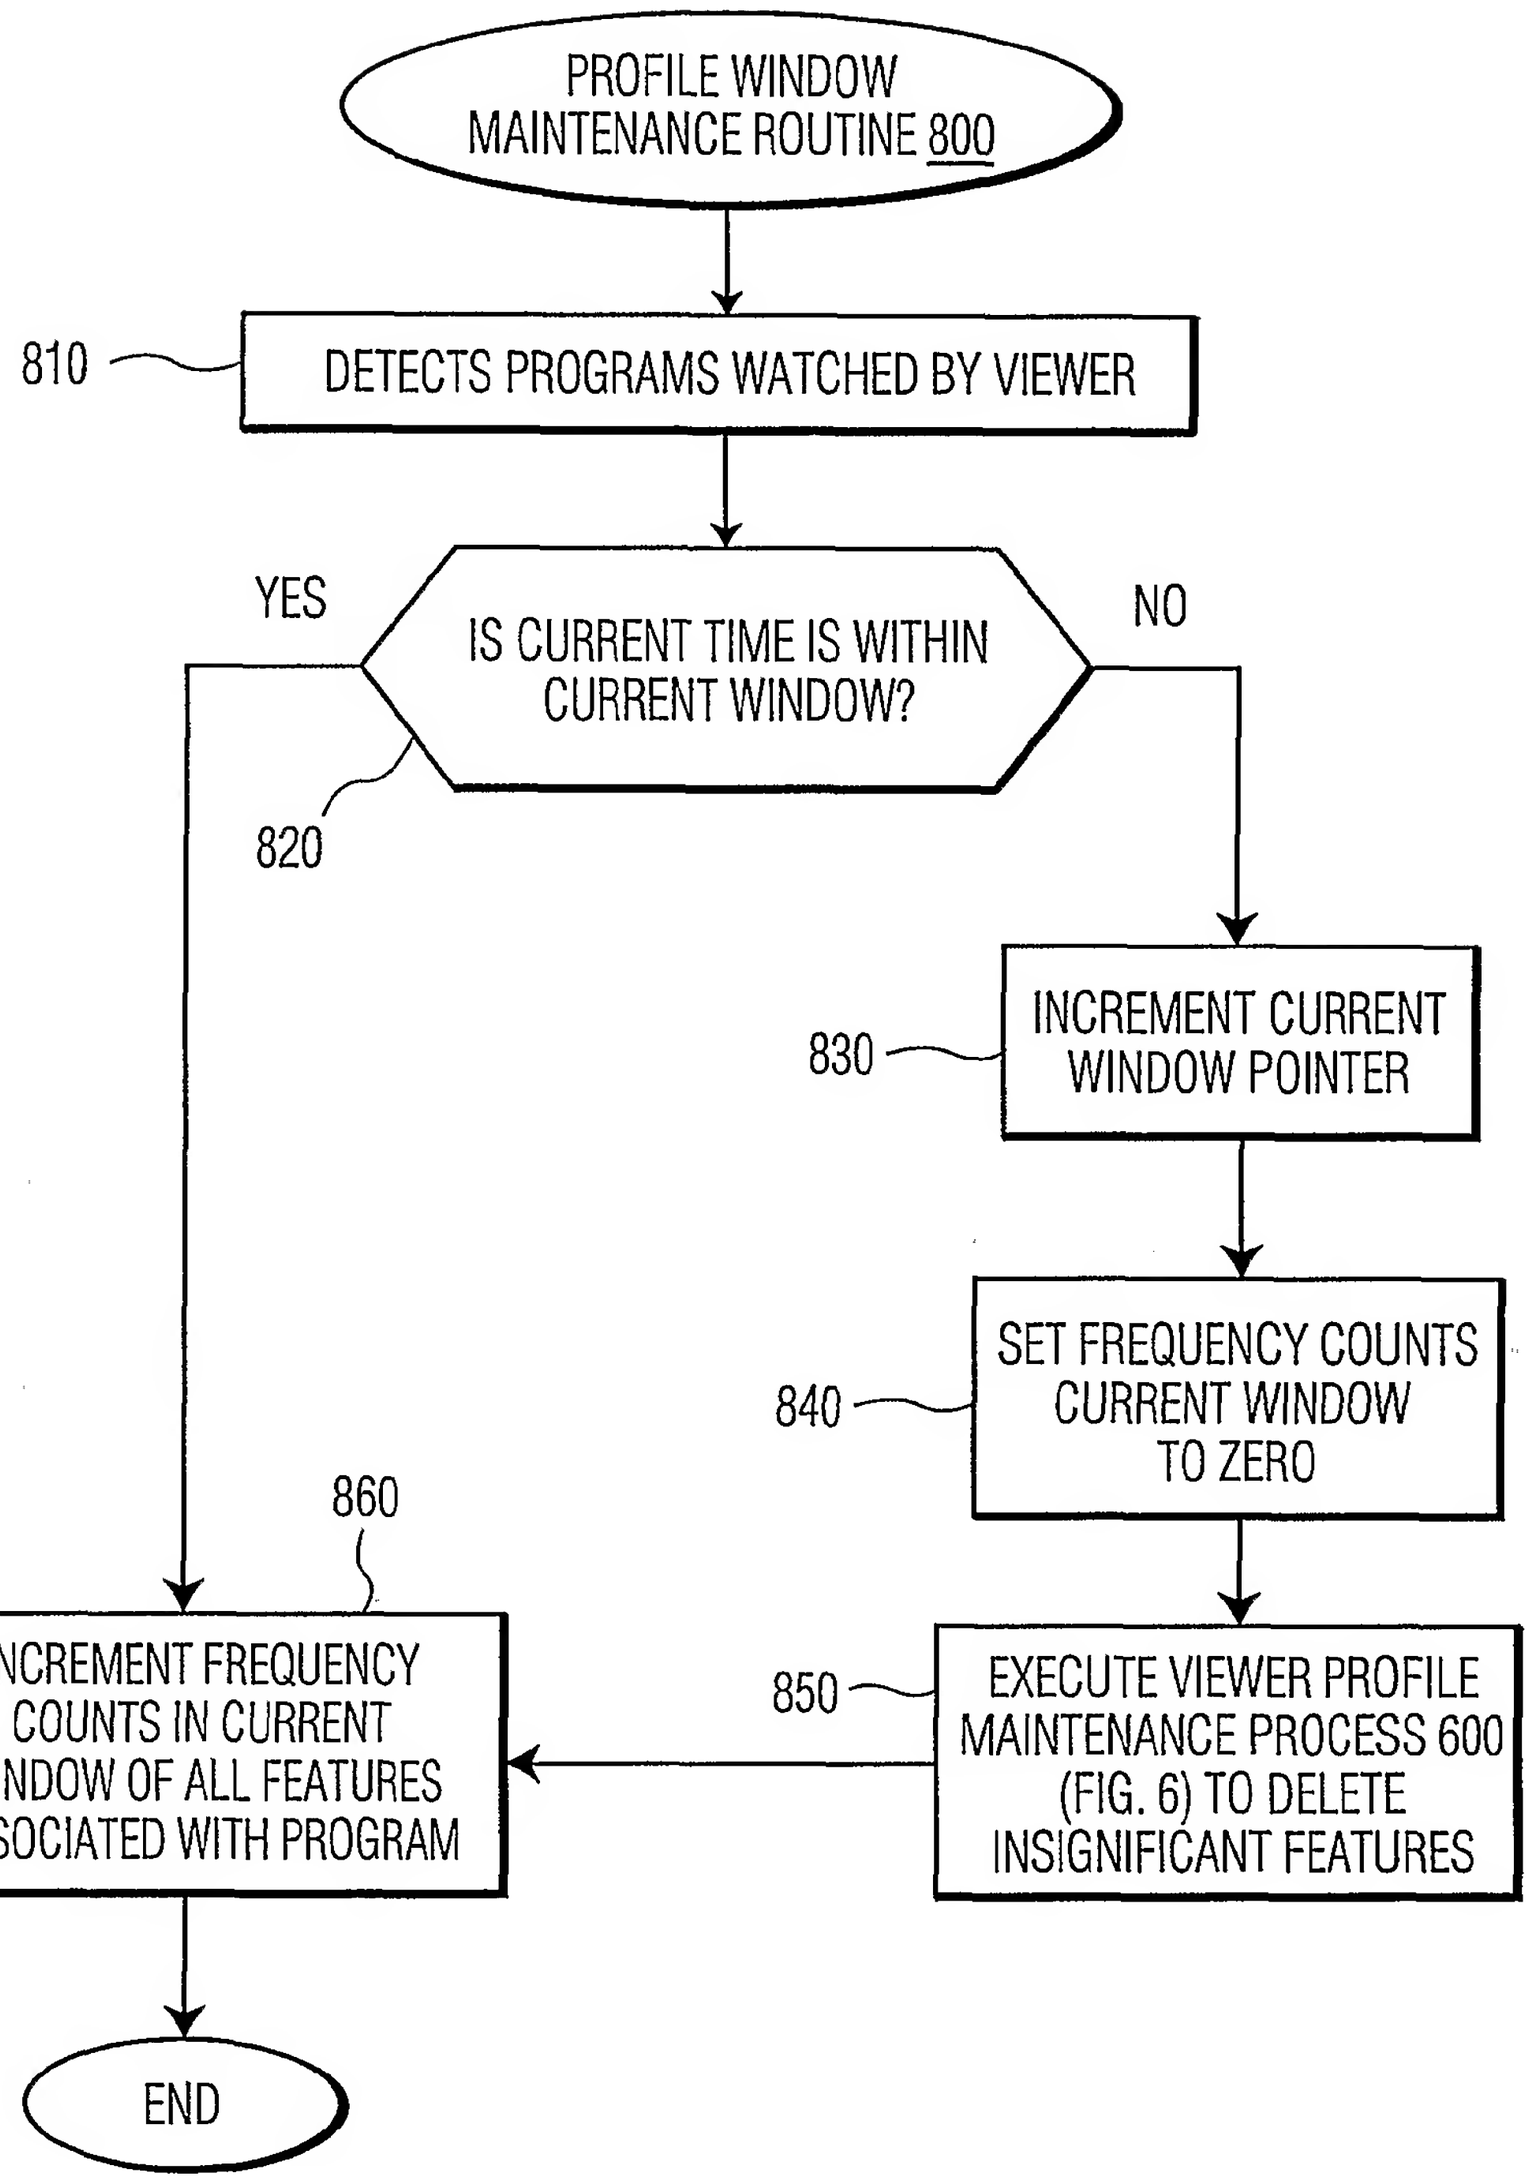
PROFILE WINDOW MAINTENANCE ROUTINE 800
810
DETECTS PROGRAMS WATCHED BY VIEWER
YES
IS CURRENT TIME IS WITHIN CURRENT WINDOW?
NO
820
830
INCREMENT CURRENT WINDOW POINTER
840
SET FREQUENCY COUNTS CURRENT WINDOW TO ZERO
860
INCREMENT FREQUENCY COUNTS IN CURRENT WINDOW OF ALL FEATURES ASSOCIATED WITH PROGRAM
850
EXECUTE VIEWER PROFILE MAINTENANCE PROCESS 600 (FIG. 6) TO DELETE INSIGNIFICANT FEATURES
END

FIG. 8